Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 15:56:15
REFERENCE NO: 270

Author Names & Affiliations

- Dan Stanzione - TACC/University of Texas at Austin

Contact Email Address (for NSF use only)

(Hidden)

Research Domain, discipline, and sub-discipline

Computational Science

Title of Submission

Future Needs for Advanced Cyberinfrastructure: A CI provider perspective

Abstract (maximum ~200 words).

The Texas Advanced Computing Center (TACC) provides cyberinfrastructure to a diverse group of researchers around the nation and world. In the past year alone, TACC supported more than 40,000 researchers at 400 institutions. TACC provided over 1 billion compute hours in support of their research, and maintains more than 60 petabytes of data in over five billion files for its users. We are well-positioned to see a cross section of CI requests across many disciplines and have noted the emergence of broad categories of users and communities. A small group of users does ground-breaking science through the largest simulations possible, and a much larger group requires large numbers of smaller simulations that together enable new science. Machine learning and data-intensive computing have brought computation to disciplines without a mathematical foundation. And, across these groups are users with different levels of computational sophistication that benefit from advanced interfaces for CI services. All of these communities now rely upon sophisticated data management (including discovery), modern and diverse computing platforms, and advanced expertise. In order to for these communities to continue to reap the benefits of advanced computation, NSF is challenged to provision CI in a way that is stable yet continuously evolving.

Question 1 Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

The Texas Advanced Computing Center (TACC) provides cyberinfrastructure services to an enormously diverse group of researchers at the University of Texas and around the nation and world. In the past year, TACC supported more than 40,000 researchers at 400 institutions with more than a million simulation and data analysis runs consuming a billion compute hours across fifteen distinct computational platforms and more than a dozen web-based portals and gateways. In addition, TACC stores and maintains more than 60 petabytes of data and more than five billion files for this community.

In our role, we are well-positioned to see a cross section of CI requests across many disciplines, which gives some perspective to the challenges researchers face and the CI constraints placed upon them. Broadly, we would note the following trends and challenges.

Discovery through Leadership-class simulation a relatively small subset of our users, but among the most published and consumers of a relatively large fraction of our cycles, push the boundaries of our largest systems to do computations at never-before-seen scales. A relatively recent example would be the work that won the 2015 Gordon Bell prize in mantle convection. This team added both new levels of fidelity to previous simulations, enhanced the numerical methods used to perform the computations, increased the level of adaptivity in their mesh, then adapted their code for new architectures (Xeon Phi on Stampede and the Blue Gene/Q at Lawrence Livermore). The result was groundbreaking – not only was the simulation (ultimately at >750,000 cores) of stunning scale, but the new insight was as well. This computation showed for the first time how the flow of the Earth's mantle creates the underlying mechanism that guides continental drift. Through the use of cyberinfrastructure at extreme scale, new and impactful science was created.

Discovery through massive throughput computing a much larger group of our researchers (and available computing) can be classed as what would commonly be called "capacity computing" – no single simulation or analysis run is individually very large, but the desired scientific results require very large numbers of these runs that collectively require significant CI resources. A spectacular recent example of this would be the 2016 LIGO experimental verification of Einstein's prediction through the first observation of gravitational waves. This achievement was hailed as one of NSF's most significant achievements by Director Cordova. The role of CI in this discovery bears closer examination: A vast and expensive experimental instrument was required. Like most modern instruments, it produces in turn a vast quantity of digital data to be analyzed. The computational requirements of analyzing this data were also, of course very large, and TACC was one of a number of centers that contributed resources to this analysis. There is a critically important subtlety here that must not go unobserved: while the computing requirements were huge and expensive, a bit of software tuning and algorithmic work made an enormous difference in just how expensive. Optimization work in conjunction with CI experts at TACC improved the performance of a key part of the computation – the Fourier Transform – by more than triple. While this still left a huge computational load to be satisfied, the reduction in the cost of throughput computation was more than $10M. When considering throughput computing issues, or perhaps any computing issues, it is important not just to consider the cycles, but how effectively they can be used, which requires expertise far beyond hardware.

Discovery through Data Intensive Computing While the LIGO example could accurately be characterized as a data intensive computing problem, the solutions could be efficiently generated on the types of computer systems and with the types of tools and techniques that are commonly used for numerical simulation. Another trend we see, with large numbers of users but currently limited scale, is in doing significantly different kinds of algorithms and data analysis to deal with data-intensive problems. Often, this involves making changes to how we configure our computing platforms (see question 2). Many new users are coming to consume CI services as a result of the broad trend of the generation of digital data – from the many new kinds of sensors and instruments (i.e. satellite images, genomic sequencers, MRI machines, Cryo-Electron Microscopy) and from personal devices connecting to networks (mobile phone GPS data, internet histories, social media postings), etc. This rise of data has created enormous computational problems in fields where traditional numerical simulation was not a driver or where numerical models don't exist. Life sciences and health have been prominently featured in this transformation, but it has spread to social sciences, arts, and media. Statistical models, Heuristics, Neural Networks, and other Machine Learning/Deep Learning techniques will frequently replace traditional differential equation based modeling. This means new software to be developed and tuned, and in many cases the underlying hardware systems can be changed to be more effective with these paradigms. A driving research example would be the DARPA-funded MEMEX project, which is combing the Deep Web to discover lurking patterns that may help identify human trafficking and other illicit trades, as well as identify national security and cyber threats. This project, led by Chris Mattmann at Cal Tech, takes advantage of the Hadoop framework on TACC's Wrangler system, which is a radical re-envisioning of the compute to I/O ration in a supercomputing platform.

Enabling Broad Discoveries through Web accessibility A final emerging trend in research we see is the migration of massive numbers of researchers who traditionally have worked at the "desktop scale" to shared CI services. Excellent examples of this in our research community would be the Cyverse project (biology and life sciences; formerly known as iPlant), the DesignSafe project (natural hazards engineering), the Galaxy project (genomics workbench), and the Science Gateways Community Institute (SGCI, common infrastructure for all the preceding ones). These projects share a number of common themes: First, they all serve large numbers of researchers; by sheer number of users, they dwarf the other types of projects – Galaxy has more than 50,000 users, Cyverse nearly 40,000, and DesignSafe attracted more than 1,000 in its first year. Each has resulted in huge numbers of impactful research publications. Each of these projects attracts users for a variety of reasons; but each reason is true in every project. For some users, it is a matter of their computational or data analysis load scaling beyond their desktop. For others, it is because the complexity of software tools and workflows and other CI resources is beyond their local skill (or time) to configure and manage, and they are looking for a simple and unified interface. For others, it is to better enable collaboration – a shared space to work with colleagues, particularly remotely, on data and results. For many others, they are driven by the need for reproducibility – to have results, and a provenance log, tracked and stored in a repository. In all of these projects, making complex or large scale resources available through the web, through a common interface, is a central element. Likewise, in all of these projects, data management in addition to CI is a common theme. They also each unite both large scale batch-oriented computing and interactive computing in a single interface, as well as uniting simulation, data analysis, and visualization. While these types of projects consume a relatively small fraction of our total computational capacity, they represent a vast majority of users, and critically important science happens at both the large and small scales.

Question 2 Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

As question 1 has shown, the researchers who come to TACC to solve their problems have diverse needs. Similarly, the CI to support them is equally diverse. In recent years, we have seen the needs for cyberinfrastructure grow in breadth, depth, and complexity. For those working in large scale simulation, the appetite for high end computing cycles is voracious and growing exponentially – we see that the demand for our compute cycles (with more than 150,000 NSF-supported processor-cores at TACC, and an additional 50,000 supported by other sources) are oversubscribed by at least a factor of five for deserving and meritorious research. But, in addition to simulation needs, we also note breadth in demand for other cyberinfrastructure services as well.

In terms of demand for computing, as we see the different research classes as described in question 1, we see a similar diversity in needs for computing systems. High Performance Computing (HPC), at the largest scale, remains in enormous and growing demand. High Throughput Computing (HTC) is equally important and equally in demand, but can be satisfied with either HPC systems or more Cloud-like systems. Cloud-style systems, both private or commercial, can work on HTC workloads, but are increasingly popular to meet other types of demands as well – for interactive computing (often in the form of Matlab, Python, Jupyter, R, Labview, or visualization tools) is a fast-growing need, and a low-utilization cloud system is an excellent way to provision interactive sessions. Likewise, a low-utilization cloud-style system is useful for meeting real-time and near real-time demands for cycles (i.e., rapid weather or flood forecasts). Our experience thus far with clouds is that while they can help meet HTC, interactive, and real-time demands, the most important characteristic for cloud adoption is they be *collocated with data sets*. We also see rising demand for specialized systems; for I/O and data intensive workloads, for machine learning/deep learning applications, and for visualization. For the most part, users don’t want to think deeply about which system they need, they simply want diverse computing options conveniently available.

There is a similar set of diverse needs in data and storage. Again, the common demand is the voracious appetite for capacity, but capacity alone is too coarse grain a story to really understand needs. There is demand for archive storage – long term, high redundancy, secure and safe, low cost, but not necessarily performant. There is a demand for what we term “collections” storage – online, accessible, results, for collaboration or publication, that need to be available much more quickly than archive, but still not with high end performance, as most consumers are on the far end of a wide area network. Then there are the higher performance tiers, of which there are appear to be at least four classes in demand; very high bandwidth, short term “scratch” storage in traditional filesystems to support HPC, other fast “working” storage for both relational and object data, and finally “high performance analytics” storage which is not only high bandwidth but also supports very high transaction rates for random access (typically implemented through flash or NVMe). Fast storage systems fall flat when computing is not nearby; we have found numerous instances where both the cost advantage and time advantage of commercial clouds disappear when storage and compute are not co-localated.

The third hardware category of CI demand we see is for networks. Again, capacity is an issue – researchers inevitably want more, and again there is some diversity – very fast datacenter-scale networks to support HPC and other tightly interconnected systems, demand for bandwidth in WAN connections, and demand for bandwidth all the way to the user’s endpoint (laptop), to support streaming visualizations and rapid data transfer. While most of these are solved with hardware investment, our experience is the “end-to-end” performance issue require significant human-in-the-loop tuning by an expert to truly realize the return on the hardware investment. One trend of concern is the rise of per-byte bandwidth charges at academic institutions – these are beginning to have a chilling effect on research projects.
Software and Algorithms can not be overlooked, but often are. New scales, new systems, and new science all inevitably require significant investment in new mathematics, algorithms and software. Software then requires effort to tune and optimize. This is perhaps the biggest gap in current CI investments – while there is substantive investment in new software capability, there has traditionally been less in maintaining and modernizing software (though there is more investment here now than several years ago), and even less investment in

tuning and performance optimization, to say nothing of verification and validation. Growing investments in this area are a must, as the returns on all other investments hinge directly on this. The LIGO example in the question 1 response is an excellent instance of this – a few months of optimization effort reduced the required computational hardware needs by millions of dollars, freeing those resources for other critical computation (and shortening the time to result for the LIGO team dramatically).

Even the combination of all of the kinds of hardware and software investments listed above are not sufficient to maximize scientific results from our CI. We can provide to our researchers excellent applications, blazing fast computing systems, and capacious storage platforms – and this simply unearths another layer of needs for data management, data curation, cybersecurity, compliance for personally identifiable data, visualization, collaboration platforms, provenance and reproducibility etc. As science becomes more cyber-centric, the scope of CI tasks in science continues to grow. The most effective method we have seen to address this – and where demand continues to explode, is for large, highly-collaborative, discipline-centric projects to provide a comprehensive, unified CI – to tie together all of the types of data and computational systems above, to provide the software, algorithmics, and performance expertise to make them work together, and to provide a set of interfaces that are familiar to the researchers in that discipline. In the successful projects, we find them not just building the (typically web) tools to create portals for these communities, but also doing two more critical things: first, providing a concentration of CI expertise where researchers can come to partner to maximize the benefit of CI in their projects. Second, building out platforms in which others in that discipline can build out their own applications and uses – not simply a static portal, but a toolkit for innovative new interfaces (this seems best enabled today through web-based application programmer interfaces).

The combination of a range of capable computing systems, abundant storage, ubiquitous fast networking, with software and algorithm expertise, good interfaces, and a deep pool of human expertise, will enable a future CI that can dramatically enhance the progress of every scientific endeavor of the NSF.

Question 3 Other considerations (maximum ~1200 words, optional): Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

We would be remiss in discussing the challenges of question 1 and the infrastructure of question 2 without also describing the human dimensions of CI-enabled research. The incredibly diverse set of research challenges, matched with the incredibly capable but equally complex CI ecosystem, is putting ever more demands on the skills of research teams to work effectively.

A significant issue we have identified is the skillset and workforce required to properly match available cyberinfrastructure capabilities to the problems we have at hand. Modern cyberinfrastructure systems can by dauntingly complex, with computers containing tens of thousands of cores, alternative accelerated architectures, and many kinds of storage to choose from. Software systems are equally complex, with simulation programs comprising hundreds of thousands of lines of code and dozens of configuration options. Other parts of the cyberinfrastructure ecosystem have equal complexity (including the options for many different providers of these services).

We observe an increasing divergence in the researchers making use of CI into two camps of “haves” and “have nots” – Sophisticated HPC users on the one hand, who employ multiple full time (usually) postdocs to deal strictly with evolving their code to dealing with the next generation of computing systems, and the users who are strictly consumers of CI who work in (for instance, perhaps) Matlab at the desktop scale. While one camp can make use of giant HPC systems with hundreds of thousands of cores and deep memory and storage hierarchies, these systems, and the programming models that go with them, are increasingly foreign to the other camp, who are focused primarily on ease of use.

There seem to be two schools of thought on proceeding – one, “train the end user”, and two “engage an expert consulting”. We believe neither are adequate. The systems of today (and 2030) are too complicated for simple training webinars and documentation to bridge the gap. Hiring a single “CI expert” can be equally inadequate, as modern CI is a diverse collection of skills rarely (never?) embodied in a single individual.

Computational science and engineering is now truly a discipline. In the future, we will increasingly need diverse cadres of CI professionals, combined with increased CI literacy among engineers and scientists in general. Centers which create a critical mass of the necessary kinds of expertise need to be broadly available for deep collaborative relationships with numerous researchers. Alternatively, or perhaps in addition, virtual organizations that allow the “computer person” in a given project to not operate on an island and share expertise with others could contribute to a solution.

In the end, our CI can not simply be purchased – we need investments in both the physical infrastructure and the people to employ it.

Consent Statement